

What is claimed:

1. A method for intelligently updating content in a server cluster having a plurality of servers to provide consistent data, comprising the steps of:

5 (a) storing pre-existing content on a server that is being updated in a temporary location;

(b) updating said server with said content;

(c) inhibiting said server from accepting requests for said content and redirecting requests for said content in said server to said temporary location;

10 (d) repeating steps (a) and (c) until each server is updated;

(e) determining if said content has been successfully updated on each server;

(f) storing said pre-existing content in a staging server and enabling said server to accept requests for said content if it is determined that said content has been
15 successfully updated; and

(h) restoring said pre-existing content to each server and enabling said server to accept requests for said pre-existing content if it is determined that said content has not been successfully updated.

2. The method of claim 1, wherein said content is a file or directory.

20 3. The method of claim 1, wherein said content is a group of files or directories.

4. The method of claim 3, wherein said content is an atomic content; and wherein the step (d) determines if said group of files or directories have been successfully updated.

25 5. The method of claim 3, wherein said content is a non-atomic content; and wherein the step (d) determines for each file or directory if said each file or directory has been successfully updated.

30 6. The method of claim 1, wherein said content represents file or directory to be removed from said server; and wherein the step (b) deletes said content from said server.

7. The method of claim 1, wherein said content represents content stored in said staging area; and wherein the step (b) restores said stored content to said server.

8. A method for intelligently updating content in a server cluster having a plurality of servers to provide consistent data, comprising the steps of:

(a) storing pre-existing content on a server that is being updated in a staging server;

(b) updating said server with said content;

(c) inhibiting said server from accepting requests for said content by a load balancer;

(d) determining if a predetermined server threshold has been met for said content;

(e) permitting said server from accepting said requests and inhibiting servers that has not been updated with said content from accepting requests if it is determined that said predetermined server threshold has been met;

(f) repeating steps (a) and (e) until each server is updated; and

(g) restoring said pre-existing content to each server and enabling said server to accept requests for said pre-existing content if it is determined that said predetermined server threshold has not been met.

9. The method of claim 8, wherein said content is a file or directory.

10. The method of claim 8, wherein said content is a group of files or directories.

11. The method of claim 10, wherein said content is an atomic content; and wherein the step (d) determines if said server threshold has been met for said group of files or directories.

12. The method of claim 10, wherein said content is a non-atomic content; and wherein the step (d) determines for each file or directory if said server threshold has been met for said each file or directory.

13. The method of claim 8, wherein said content represents file or directory to be removed from said server; and wherein the step (b) deletes said content from said server.

14. The method of claim 8, wherein said content represents content stored in said staging area; and wherein the step (b) restores said stored content to said server.

15. An intelligent content distributor for intelligently updating content in a server cluster having a plurality of servers to provide consistent data, comprising:

a console for generating a job for updating said cluster with said content;
a scheduling for scheduling said job; and

an executor for executing said job for each server in said server cluster, wherein said job comprises: storing pre-existing content on a server that is being updated in a temporary location, updating said server with said content, inhibiting said server from accepting requests for said content and redirecting requests for said content in said server to said temporary location, and determining if said content has been successfully updated on said server; and

wherein said executor is operable to store said pre-existing content in said intelligent content distributor and enabling said plurality of servers to accept requests for said content if it is determined that said content has been successfully updated; and

wherein said executor is operable to restore said pre-existing content to each server and enabling said plurality of servers to accept requests for said pre-existing content if it is determined that said content has not been successfully updated.

16. The intelligent content distributor of claim 15, wherein said job comprises publishing, deleting or restoring content.

17. The intelligent content distributor of claim 15, wherein said scheduler is operable to reschedule said job if it is determined that said content has not been successfully updated.

18. An intelligent content distributor for intelligently updating content in a server cluster having a plurality of servers to provide consistent data, comprising:

a console for generating a job for updating said cluster with said content;
a scheduler for scheduling said job;

an executor for executing said job for each server in said server cluster, wherein said job comprises: storing pre-existing content on a server that is being updated in said intelligent content distributor, updating said server with said content,

and determining if a predetermined server threshold has been met for said content;
and

a load balancer for inhibiting a server that has been updated with said content
and from accepting requests until it is determined that said predetermined threshold
5 has been met; and

wherein said executor is operable to restore said pre-existing content to each
server and enabling said plurality of servers to accept requests for said pre-existing
content if it is determined that said predetermined threshold has not been met.

19. The intelligent content distributor of claim 15, wherein said job
10 comprises publishing, deleting or restoring content.

20. The intelligent content distributor of claim 15, wherein said scheduler
is operable to reschedule said job if it is determined that said predetermined threshold
has not been met.